## REMARKS/ARGUMENTS

These remarks are made in response to the Final Office Action of September 20, 2004 (Office Action). As this response is timely filed within the three-month statutory period, no fee is believed due.

In paragraphs 3 and 4 of the Office Action, Claims 1-3, and 15 have been rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent Number 6,061,653 to Fisher, *et al.* (Fisher). In paragraphs 5-6, Claim 4 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Fisher. In paragraph 7, Claims 5-10, 12-14, and 17-21 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Fisher in view of U.S. Patent Number 6,539,087 to Walsh, *et al.* (Walsh). In paragraph 8, Claims 11 and 16 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Fisher in view of U.S. Patent Number 6,535,513 to Kao, *et al.* (Kao).

It may be helpful to briefly review the features of Applicants' invention before addressing the rejections on the art. The Applicants' invention is directed to a speech processing board that is used in high volume speech processing applications. The speech processing board can be deployed both in a conventional computer telephony (CT) architecture and in a voice over IP (VoIP) gateway/endpoint architecture. The speech processing board of the present invention also can accommodate multiple instances of text-to-speech (TTS) application tasks and small vocabulary speech recognition tasks.

Unlike conventional speech processing boards, Applicants' speech processing board includes multiple processor modules, each of which can execute multiple instances of full function, large vocabulary speech recognition tasks similar to those of a conventional speech recognition engine with shared memory. Moreover, Applicants' speech processing board further includes a language model cache mapped to a common address space such that the language model cache can be uniformly accessed by each of the multiple processor modules. Similarly, a method aspect of the present invention includes loading selected language models in a storage separate from the multiple

processor modules such that the selected language models are uniformly accessed by each processor module at a common address.

At pages 2-3 of the Office Action, in rejecting independent Claim 1, the Examiner asserts that Fisher's disclosure of a process-independent, model storage area is the "functional equivalent" of a uniformly accessible language model cache. The specific portion of Fisher cited by the Examiner includes the following description:

> "data in process-independent storage area **42a** may be shared by more than one speech recognition process and, as such, it need not be duplicated for different processes and therefore the amount of overall memory space is reduced." (Col. 6, lines 48-52.)

What is important to recognize in reading this language, in the context of Fisher in its entirety, is that it is the process-independent data itself, not a language model cache mapped to a common address that is shared. This point is made throughout the other portions of Fisher. For example, elsewhere Fisher states that

> "because this data [pertaining to the process-independent portion of a speech model] may be shared by certain processes, it is possible that such data can be loaded during initialization and then not re-laded each time a new process which may share that data is performed." (Col. 6, lines 53-56.)

What is crucial is the question of where is this data loaded during initialization. Fisher makes clear that the data is not loaded to a cache mapped to a common address as with Applicants' invention. Instead, when the process-independent data is to be used in processing a call, it is, as Fisher explicitly states

> "loaded from hard storage into the memory associated with a DSP **18** either at boot time or some other infrequent interval." (Col. 7, lines 49-53.)

Since, as FIG. 1 of Fisher clearly illustrates, there are multiple DSPs, it follows that in Fisher it is the process-independent data that is shared with each of the DSPs, not a language model cache mapped to a common address. This is much more than a difference of semantics. When the shared data in Fisher is to be used by a particular one of the multiple DSPs, it is transferred from hard storage to the memory associated with DSP. (Col. 7, lines 49-53.) Fisher provides no language model cache, and there is none that is mapped to a common address that is uniformly accessible by each DSP as with Applicants' invention.

This difference between Fisher and Applicants' invention is made even more stark when one compares how data pertaining to the process-independent portion of a speech model is shared with the DSPs in Fisher. Fisher relies on a "host processor" **16** for managing calls and pushing the data pertaining to the process-independent portion of a speech model, as well as that pertaining to the process-dependent portion, down to individual DSPs **18** over a common bus. (See FIG. 1) With Fisher, the data pertaining to the process-independent portion of a speech model are pushed down during initialization to a memory associated with an individual DSP. The data is then run individually by the particular DSP in Fisher. It is not shared in Fisher among or between DSPs.

Moreover, when data pertaining to the process-dependent portion is pushed down as needed, Fisher must flag the data to invalidate previous process-dependent language model data. This further illustrates that the DSPs are connected via a common I/O bus, but that Fisher lacks the memory coherence provided by Applicants' invention. This lack of memory coherence stems from Fisher's not providing a language model cache that is mapped to a common address and that is uniformly accessible by each of different processor modules.

Applicants' invention provides enhanced performance over the push/invalidate flag model of Fisher. Applicants' invention can function independently of a host processor in pushing data to its working elements, namely, the multiple processor modules. The data

can be pulled on demand from the language model cache in Applicants' invention. Moreover, the mapping of the language model into a common memory address range, permits a leveraging of the normal caching hardware built into processors such as RISC processors. These and other significant advantages follow from the differences in Applicants' invention over Fisher. For example, with Applicants' invention updates of the language model cache locations can be "snooped" by a processor and their local copies subsequently invalidated.

For the various reasons already stated, Fisher fails to teach each feature recited in independent Claim 1. Therefore, Applicants respectfully request withdrawal of the 35 U.S.C. § 102(e) rejection with respect to independent Claim 1. Whereas dependent Claims 2, 3, and 15 each depend from Claim 1 and recite additional features, Applicant's respectfully request that the rejection regarding these claims similarly be withdrawn.

With regard to the Examiner's rejection of claims 5-10, 12-14, and 17-21 under 35 U.S.C. § 103(a) in view of Fisher and Walsh, and of claims 11 and 16 under 35 U.S.C. § 103(a) in view of Fisher and Kao, Applicants reiterate that, regardless of whether the references are properly combinable, the combination of Fisher with Walsh, as well as that of Fisher with Kao, fails to teach every feature recited in the claims. For example, none of the references disclose expressly or implicitly a language model cache that communicatively links to a bridge and that is uniformly accessible by each of multiple processor modules at a common address, as recited in amended Claim 17.

Nor do any of the references teach or suggest a selected language model that is uniformly accessible in a common address space by each of the multiple processor modules as recited in independent Claims 18 and 20, as amended. Neither Walsh nor Kao is cited as teaching or suggesting such a feature, and, as discussed above, Fisher fails to even suggest such a feature.

Therefore, since the references singly and jointly fail to teach every feature of the claimed invention, the Applicants respectfully request that the 35 U.S.C. § 103(a)

rejection with respect to independent Claims 17, 18, and 20 be withdrawn. Whereas the remaining dependent claims each recite additional features, withdrawal of the 35 U.S.C. § 103(a) rejection of these claims is also respectfully requested.

Applicants believe that this application is now in full condition for allowance, which action is respectfully requested. Applicants request that the Examiner call the undersigned if clarification is needed on any matter within this Amendment, or if the Examiner believes a telephone interview would expedite the prosecution of the subject application to completion.

Respectfully submitted,

Date: 11/22/04

Gregory A. Nelson, Registration No. 30,577
Richard A. Hinson, Registration No. 47,652
AKERMAN SENTERFITT
Customer No. 40987
Post Office Box 3188
West Palm Beach, FL 33402-3188
Telephone: (561) 653-5000